

Spatial versus Object Working Memory: PET Investigations

Edward E. Smith, John Jonides, Robert A. Koeppe, Edward Awh, Eric H. Schumacher, and Satoshi Minoshima

University of Michigan

Abstract

■ We used positron emission tomography (PET) to answer the following question: Is working memory a unitary storage system, or does it instead include different storage buffers for different kinds of information? In Experiment 1, PET measures were taken while subjects engaged in either a spatial-memory task (retain the position of three dots for 3 sec) or an object-memory task (retain the identity of two objects for 3 sec). The results manifested a striking double dissociation, as the spatial task activated only right-hemisphere regions, whereas the object task activated primarily left-hemisphere regions. The spatial (right-hemisphere) regions included occipital, parietal, and prefrontal areas, while the object (left-hemisphere) regions included inferotemporal and parietal areas. Experiment 2 was similar to Experiment 1 except that the stimuli and trial events were identical for the spatial and object tasks; whether spatial

or object memory was required was manipulated by instructions. The PET results once more showed a double dissociation, as the spatial task activated primarily right-hemisphere regions (again including occipital, parietal and prefrontal areas), whereas the object task activated only left-hemisphere regions (again including inferotemporal and parietal areas). Experiment 3 was a strictly behavioral study, which produced another double dissociation. It used the same tasks as Experiment 2, and showed that a variation in spatial similarity affected performance in the spatial but not the object task, whereas a variation in shape similarity affected performance in the object but not the spatial task. Taken together, the results of the three experiments clearly imply that different working-memory buffers are used for storing spatial and object information. ■

INTRODUCTION

The construct of working memory is central to our understanding of higher-level cognition. The few broad-based theories of cognition that exist—notably the computational models of Anderson (1983) and Newell (1990)—include working memory as a critical component of the cognitive architecture, and give it a major role in reasoning, problem solving, and language understanding. Other computational models also assume that the fundamental limit on solving a problem or understanding a complex sentence is the limited capacity of working memory (e.g., Carpenter, Just, & Shell, 1990; Just & Carpenter, 1992). In addition, cognitive research that does not rely on computational modeling supports the view that the limits and properties of working memory are manifested in a myriad of psychological tasks (see Jonides, 1995, for a recent review). Despite the agreed-upon centrality of working memory, the most fundamental questions about its nature remain unanswered. In this report, we bring to bear the methodology of functional brain imaging to answer a basic question about working memory.

The general issue at stake is whether working memory

is unitary in character, or whether qualitatively different working memories are employed for different kinds of information. Baddeley and his colleagues (e.g., 1986, 1992) have been leaders in addressing this issue, and their research implicates at least two different working memories, one that holds verbal material in a speech-like code, and another that holds visuospatial material in some sort of pictorial code. A second dissociation that has been considered is that between working memories for visual and spatial information (e.g., Baddeley & Lieberman, 1980). However, behavioral research on this issue has been indecisive. This latter issue is the focus of the present paper. Whereas previous studies on the distinction between visual and spatial working memories have relied exclusively on strictly behavioral experiments, we combine behavioral experiments and studies using positron emission tomography (PET) measures of regional cerebral blood flow to provide evidence for separate visual and spatial working memories.

Perhaps the most compelling strategy for demonstrating different kinds of memories is the logic of double dissociation. In strictly behavioral studies, this logic takes the familiar form:

If there is a behavioral factor that influences performance on Task A but not Task B, and another behavioral factor that influences performance on Task B but not Task A, then these two tasks are mediated by different processing mechanisms.

To extend this logic to measures of neural activation or cerebral blood flow, one assumes that different processing mechanisms have different neuroanatomical loci. Then the logic becomes:

If performance on Task A is associated with heightened neural activity in Brain Region a but not Brain Region b, whereas performance on Task B is associated with heightened neural activity in Region b but not Region a, then the two tasks are mediated by different processing mechanisms.

In addition to its applicability to PET measures, double-dissociation logic applies straightforwardly to single-cell recordings with nonhuman primates. Recently, Wilson, O'Scalaidhe, and Goldman-Rakic (1993) applied the logic to the issue of visual vs. spatial working memory. In a task that has previously been shown to involve spatial working memory (e.g., Goldman-Rakic, 1987), adult monkeys were trained to remember the spatial location of a single target stimulus for a 2.5-sec delay period before executing a saccade to that location. Neurons in the principal sulcus region of dorsolateral, prefrontal cortex were found to be active only during the delay period, when the animal was presumably storing the position of the previously presented target. By contrast, in an object task used by Wilson et al. (1993), monkeys were trained to remember a nonspatial visual object, such as a simple pattern or a monkey's face, for a 2.5-sec delay period. In this case, neural activity during the delay period was revealed only in neurons in the inferior convexity region of the dorsolateral prefrontal cortex, a region directly inferior to the principal sulcus. This pattern of results strongly implies that principal-sulcus mechanisms mediate (at least in part) spatial working memory, whereas inferior-convexity mechanisms mediate (in part) object working memory.

Functional neuroimaging techniques allow one to obtain comparable data in humans on the dissociation of visual and spatial memory. One advantage of these techniques is that they can reveal activation in the entire brain, not just in one area. Accordingly, in our experiments we used PET to measure cerebral blood flow while subjects engaged in either spatial or object working-memory tasks, in an attempt to determine regions that are either more active during the spatial task or more active during the object task. We had two goals. One was to determine whether PET measures manifest a double dissociation between spatial and object working-memory tasks. A second goal was to use the PET measures to begin to determine the specific brain circuits involved in spatial and object working memory.

EXPERIMENT 1

Rationale and Tasks

In Experiment 1, different groups of subjects performed either a spatial or an object working-memory task while PET measures were recorded. Because some of the results for these tasks have been reported elsewhere (Jonides, Smith, Koeppe, Awh, Minoshima, & Mintun, 1993; Smith & Jonides, 1994), here we emphasize the contrast in results between the spatial and object tasks.

The spatial-memory task was based in part on prior work of Goldman-Rakic and her colleagues (e.g., Goldman-Rakic, 1987; Funahashi, Bruce, & Goldman-Rakic, 1989). The task is presented schematically at the top of Figure 1. On each trial, subjects began by fixating a cross in the center of a screen for 500 msec. The cross was replaced by three dots randomly arrayed on the circumference of an imaginary circle that was centered on the cross. The dots remained in view for 200 msec—too brief a duration for a successful saccade to the dots, on average. Next, the fixation cross appeared alone for a retention interval of 3000 msec. The retention interval was followed by a location probe, which consisted of a single outline circle that either encircled the location of one of the previous dots (with probability 0.50) or not. The probe circle was either centered directly over the location of a previously presented dot, or it missed the nearest dot location by at least 15 degrees. The subjects' task was to press a response button once or twice, within a 1500-msec interval, to indicate whether or not the probe encircled the location of a target dot (we use the term "target" to refer to an item that has to be remembered).

Successful performance on this spatial-memory task requires numerous processes in addition to storing spatial information in working memory, e.g., attending to the inputs, encoding spatial information, selecting and executing a response. As in the single-cell recording work, our goal is to focus on the storage component. To remove the effects of these nonstorage processes from the activation patterns of the PET images, a control condition was devised that included these processes but not storage. (This follows the subtraction methodology developed by Posner, Peterson, Fox, & Raichle, 1988.) This spatial-control task is schematized at the bottom of Figure 1. Again each trial began with a fixation cross, but in this condition the cross alone remained in view for 3500 msec (the duration of the fixation plus retention intervals in the spatial-memory task). The three dots were then presented for 200 msec, followed immediately by an interval of 1500 msec during which the three dots and probe circle were presented simultaneously. As in the memory condition, subjects pressed a response button once or twice to indicate whether or not the probe encircled a dot, but in this case their response was based on the presentation of probe and dots simultaneously so that no memory of dot location was needed. Thus this

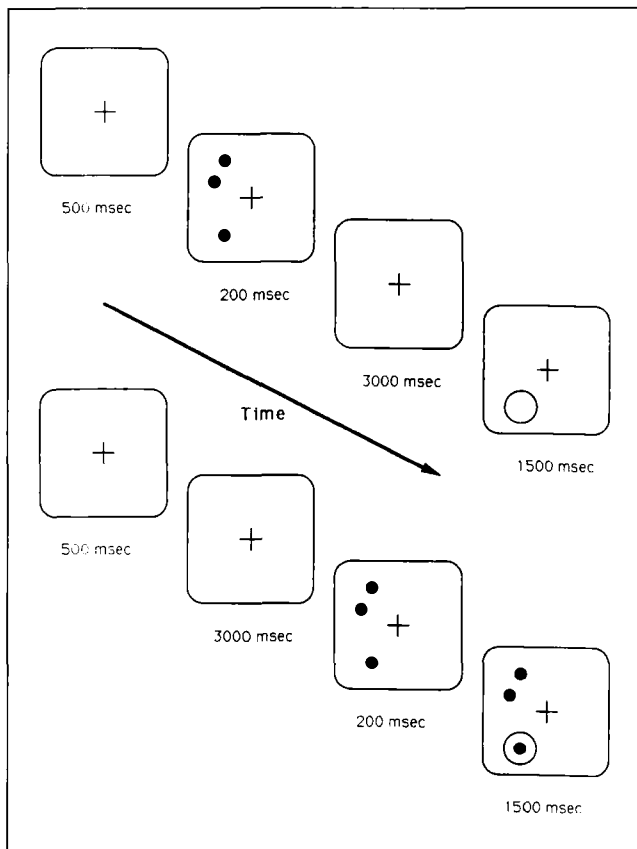


Figure 1. Schematic representation of a trial in the spatial-memory task (top) and in the spatial-control task (bottom) in Experiment 1.

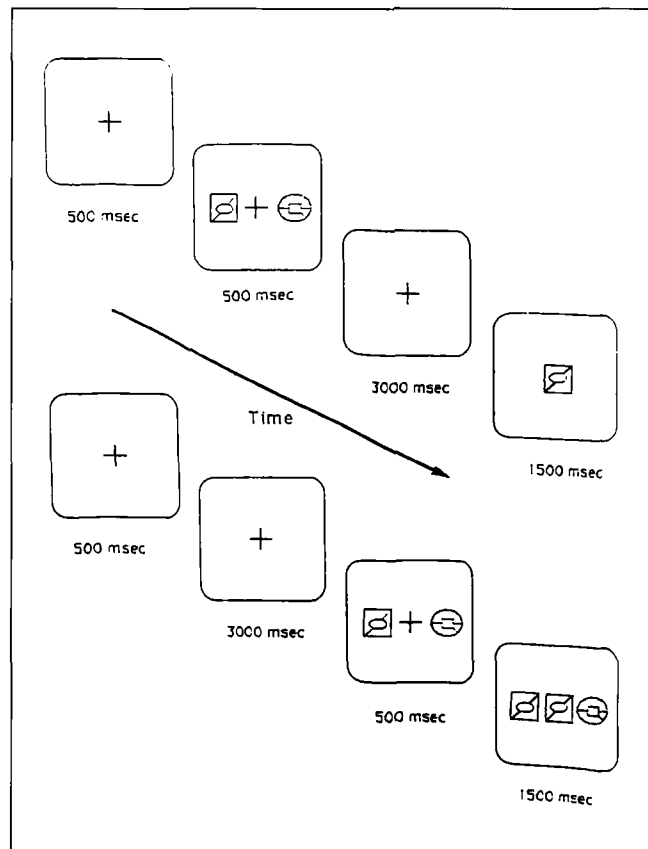


Figure 2. Schematic representation of a trial in the object-memory task (top) and in the object-control task (bottom) in Experiment 1.

spatial-control task includes very similar trial-events to the spatial-memory task, but no storage requirement.

The object-memory task and its control condition are schematized in Figure 2. The object-memory task, illustrated at the top of Figure 2, has a similar structure to the spatial task. Each trial began with a fixation cross exposed for 500 msec. The cross alone was replaced by two objects presented on either side of the cross. These remained in view for 500 msec (the minimum duration needed for accurate encoding of the objects, as determined by pilot work). The objects were unfamiliar geometric figures, constructed so that each consisted of an outline shape with a second shape embedded within it, with lines at various orientations connecting the inner and outer shapes. The third trial event was a return to the fixation cross alone, for a retention interval of 3000 msec. This was followed by a probe that consisted of an object located where the fixation cross had been. The probe was either identical to one of the two target objects (with probability 0.50) or not. The subjects' task was to press a button once or twice, within a 1500-msec interval, to indicate whether or not the probe object was identical to one of the target objects. As in the spatial-memory task, there was need for a control condition that included all but the storage processes in the object-memory task. This object-control task is schematized at the

bottom of Figure 2. It involves approximately the same trial events as the memory task, but with the events rearranged so that successful performance does not require any storage of object information. Subjects responded with one or two key presses to indicate whether or not the center object in the last display matched either of the two flanking objects.

One group of 18 subjects was given 2 blocks of trials of each of the spatial-memory and spatial-control tasks. Another group of 12 subjects received 3 blocks of trials of each of the object-memory and object-control tasks. All subjects were right-handed adults. They were given practice on their respective tasks prior to the experiment proper.

Results and Discussion

Behavioral Results

Performance on the control tasks should have exceeded that in the memory conditions because the former includes no storage requirements. This expectation was borne out. In the spatial conditions, accuracy was 100% in the control task versus 84% in the memory task; in the object conditions, performance equaled 94% in the control task versus 80% in the memory condition. Averaging over spatial and object conditions, performance

Table 1. Stereotactic Coordinates, *t* values, and Corresponding Brodmann Areas for Regions Showing Significant Activation in Experiment 1 (the Coordinates Are for the Peak Activation in the Region)

Task	Stereotactic Coordinates			<i>t</i> Statistic	Brain Area
	<i>x</i>	<i>y</i>	<i>z</i>		
Spatial	-33	19	-2	5.13	Right 47 (inferior dorsolateral frontal)
	-28	-71	29	5.51	Right 19 (occipital)
	-37	-42	38	5.92	Right 40 (posterior parietal)
	-28	-1	45	5.00	Right 6 (premotor)
Object	48	-58	-11	5.28	Left 37 (inferior temporal)
	39	3	29	4.43	Left 44 (inferior frontal)
	35	-42	34	4.29	Left 40 (posterior parietal)
	-1	14	43	4.58	32 (anterior cingulate)

was significantly better in the control than the memory tasks [$t(29) = 7.06, p < 0.001$].

A second and more diagnostic finding concerns the spatial-memory task, specifically performance on *distractor* trials in this task (i.e., trials in which the probe did not encircle a dot location). We divided the distractor trials in half on the basis of whether the probe was relatively near a dot location (between 15° and 25°) or relatively far away (between 40° and 50°). If subjects stored information that is truly spatial, performance should have been better on far than near distractors (because less precise spatial information is needed to reject a far distractor). This was the case, as subjects correctly rejected 92% of the far distractors but only 71% of the near distractors [$t(17) = 3.56, p < 0.01$]. There was no comparable variation of object similarity in the object task.

PET Results


For both the spatial and object conditions, difference images were created by subtracting activation in the control task from that in the memory task. These difference images were then analyzed by performing post hoc *t* tests on a voxel-by-voxel basis and correcting the outcomes for multiple comparisons (see the Methods section at the end of the article). This analysis results in a map of cerebral areas in which there were significant differences in regional cerebral blood flow between memory and control tasks.

The top half of Table 1 gives the results for the spatial tasks. Each row gives the stereotactic coordinates (Talairach & Tournoux, 1988) of a region showing significant activation, the *t* value obtained, and the corresponding Brodmann Area. The bottom half of Table 1 provides the comparable results for the object tasks. These same results are presented pictorially in Figure 3. The top half of the figure contains four brain images showing the significant areas of activation in the spatial task; the

bottom half of the figure contains the images for the significant areas in the object tasks. The arrowhead beside each image marks the site of activation that is of interest, while the number below gives the image's axial or *z* coordinate (an indication of how inferior or superior the relevant brain region is). These areas of activation have been superimposed on a magnetic resonance image so as to provide better anatomical localization.

Starting with the results for the spatial tasks, note first that all four regions of activation are in the right hemisphere. The most inferior of these regions is in dorsolateral prefrontal cortex, a region inferior to that implicated in single-cell studies of working memory (e.g., Wilson et al., 1993). The activity of current interest is centered in Area 47 ("Area" is used here as a shorthand for "Brodmann Area"), which may not be homologous to the site of neuronal activity in the single-cell studies. Moving up the brain, the next two sites of activation are Area 19 in occipital cortex and Area 40 in parietal cortex. Finding activation in these two regions is consistent with prior research on human visuospatial processing and memory. Area 19 has been found to be active when subjects are explicitly instructed to maintain visual images of simple forms (Kosslyn, Alpert, Thompson, Maljkovic, Weise, Chabris, Hamilton, Rauch, & Buonanno, 1993), and it is reasonable to suppose that our subjects had to maintain visual images of the dot positions. Area 40 is well known for its involvement in spatial processing and memory, as lesions in this region frequently lead to impairments in spatial coding and memory for spatial information (e.g., Warrington & Rabin, 1970; Egly, Driver, & Rafal, 1994). The final region of activation in the spatial task is Area 6 in premotor cortex. It is something of an enigma. On first thought, it suggests a motoric component, but subjects always made their responses with their right hands which should have led to activation in left-hemisphere motor regions: besides, any motor component should have been required equally in the control task and hence subtracted out of the difference image.¹

Figure 3. PET images of the four significant activation sites in the spatial-memory task (top), and of the four significant activation sites in the object-memory task (bottom). Each image is superimposed on an MRI image of a composite brain. This composite was created by first selecting the MRIs of five subjects not in the present experiment, then reorienting each one in stereotaxic coordinates, and finally averaging the reoriented MRIs. The top-right and bottom-left images in the spatial task and the bottom-left image in the object task are cut off due to the fact that the remaining activation fell outside the field of view of the camera. (The images were selected to show the extent of the activations.)



The last spatial-task result that deserves mention is a negative one. Activation in Area 46 of dorsolateral prefrontal cortex did not approach significance, yet this area may be the closest homologue to the principal sulcus region in monkey where Goldman-Rakic and her co-workers consistently find neurons that are active in spatial working-memory tasks (e.g., Wilson et al., 1993). Not

too much should be made of this negative result, though, since the next experiment will produce some support for Area 46's involvement in human spatial working memory in a different task.

Turning now to the results for the object-task (bottom half of Table 1 and Fig. 3), note first that three of the regions of activation are in the left hemisphere while the

fourth region is in a mid-line structure. This pattern, coupled with that for the spatial task, establishes a clear-cut double dissociation between the object and spatial tasks. Moving up the brain, the regions involved are Area 37 in inferotemporal cortex, Area 6 in premotor cortex, Area 40 in parietal cortex, and Area 32, the anterior cingulate. The inferotemporal and parietal activation sites are compatible with prior findings on object processing and memory. Inferotemporal cortex is known to be activated in normal object recognition (see, e.g., Grady, Haxby, Horwitz, Ungerleider, Shapiro, Carson, Herscovitch, Mishkin, & Rapoport, 1991), and lesions in this region, particularly in the left hemisphere, are associated with deficits in object recognition (Farah, 1991). Furthermore, lesions in Area 40 are associated with impairments of working memory (e.g., Kolb & Whishaw, 1990; Shallice, 1988). Activation in the anterior cingulate may also permit a straightforward interpretation. This region has been activated in prior PET studies that required selective attention and/or substantial demands on processing capacity (e.g., Pardo, Pardo, Janer, & Raichle, 1990), and the current object task seems to have had these requirements even more than our spatial task, e.g., behavioral accuracy was lower in the object task. (Even in the spatial task, though, there was detectable activation in the anterior cingulate; its level just failed to reach significance).

The final region involved in the object tasks is part of the premotor cortex, just superior to Area 44, Broca's area, known to be involved in speech. This activation site raises the possibility that subjects sometimes verbally described the objects to themselves and then silently rehearsed these descriptions during the retention interval. This kind of verbal coding was not anticipated in this study—our concern was with *visual* memory for objects not *verbal* memory for their descriptions—and suggests that our object working memory task was not completely successful in excluding verbal processes.²

At a gross level of analysis, our activation results manifest a striking double dissociation between spatial and object tasks, with spatial tasks leading to activation in right-hemisphere regions and object tasks leading to activation primarily in left-hemisphere regions. This double dissociation rules out the possibility that the additional areas active in the memory tasks compared to the control tasks reflect only increased effort or attention, for this kind of hypothesis cannot account for why *different* areas are recruited for spatial- and object- memory tasks. Rather, the double dissociation strongly implies that different systems mediate spatial and object working memory.

At a more fine-grained level of analysis, the known functionality of the particular regions activated suggest something about the neural circuits involved. One possibility is to assume that many of the activated regions implement routine storage functions. Thus the spatial-memory system might involve occipital mechanisms that

aid in the generation of (imaginal) memory representation, along with parietal and prefrontal mechanisms that are involved in aspects of storage; the object-memory system might involve inferotemporal mechanisms that generate the representation of the target forms, along with parietal mechanisms that are involved in aspects of storage. In this view, the processes that occur during the storage interval could be relatively *passive* ones, e.g., heightened neural activity that gradually diminishes with time. An alternative proposal is that some of the activated regions reflect *active* maintenance processes that occur during the retention interval, i.e., processes that "rehearse" visual representations (analogous to how inner speech rehearses verbal representations). Under this rehearsal view, the spatial- memory system might be configured as follows: The prefrontal and premotor areas, which are relatively close to motor areas, might mediate covert shifts of attention from one target dot position to the next (spatial rehearsal), whereas the parietal and occipital regions would implement more passive aspects of storage. There is some independent evidence for the proposed rehearsal component: Alivisatos and Milner (1989) found that patients with frontal lesions were impaired in a selective-attention task, and attributed this deficit to a difficulty in directing covert attention voluntarily to a particular spatial location. In the object-memory system, the inferotemporal region would mediate rehearsal or reprocessing of the form information, whereas the parietal region would implement more passive aspects of storage.

There are reasons, however, to be cautious about interpreting the active regions as mediating either passive or active visual-spatial storage. Four such reasons are described in what follows.

EXPERIMENT 2

Rationale and Tasks

There are four limitations of Experiment 1's tasks that could compromise our interpretation of the specific regions involved:

Describable Objects

As noted, some of the items used in our object task may have permitted verbal descriptions. As a consequence, the results for the object task may partially reflect *verbal* working memory. The activation in the left premotor area may reflect verbal rehearsal of an object description, whereas the activation in left parietal cortex may reflect the storage of a verbal representation rather than a visual-object one.

Different Stimuli

Although our spatial and object tasks had virtually the identical structure, the tasks used different stimuli—dots

vs. geometric objects. This raises the possibility that the differential activation patterns in the two tasks partly reflect specific stimulus effects. For example, whereas our objects contained internal details, our dots were homogeneous circles, and perhaps this difference contributed to our differential results for the object and spatial tasks. The difference in stimuli was also correlated with a difference in attentional demands. In the spatial tasks, subjects had to attend to location and ignore the identity of the stimulus, which should have been relatively easy given that there was only one kind of object—a dot—and a simple one at that. In the object tasks, subjects had to attend to the identity of the stimuli and ignore their location, which might have been more difficult because there were two locations to ignore.

Nondemanding Controls

In both our spatial and object control conditions, subjects could base their response on just their perception of the probe event (e.g., respond “yes” only if the probe circle encloses a dot). This had the desired effect of eliminating any storage component, but it may also have permitted subjects to avoid encoding the location of the target dots into memory. Consequently, our activation results may reflect primarily encoding, not storage, e.g., the occipital and parietal activation in the spatial tasks may be due to encoding dot locations, not storing them. (Our results also may reflect retrieval processes, because the control condition did not require any memory retrieval either.)

Noncommensurate Controls

In both our spatial and object control conditions, the final event in each trial consisted of the probe *plus* the targets. Such a display differs substantially from that of the final trial event in the memory conditions, which consisted of just the probe. This difference in visual events between the memory and control conditions may have promoted unwanted encoding differences between the conditions.

All four of these problems can be minimized by a change in task, schematized in Figure 4. The sequence of trial events in this paradigm is *identical* for the spatial and object tasks—only the instructions differ (see below). The memory version of both tasks unfolded as follows (see top of Fig. 4). Each trial again began with a fixation cross in the center of the screen for 500 msec. The cross was then replaced by two objects, which were irregular polygons that have been used in previous research and have been shown to be difficult to name (Attneave & Arnoult, 1956; Vanderplas & Garvin, 1959). The two objects were randomly arranged on the circumference of an imaginary circle that was centered on the fixation cross, and were exposed for 500 msec. Next, the fixation cross appeared alone for a retention interval of

3000 msec. The retention interval was followed by a probe, presented for 1500 msec, which consisted of a single object, an irregular polyhedron. The decision that the subjects had to make about the probe depended on whether they were performing a spatial- or an object-memory task. In the spatial task, subjects had to decide whether or not the probe object was in the same position as either of the target objects; in the object task, subjects had to decide whether or not the probe object was identical to either target object. In both conditions, subjects signaled an affirmative response by pressing a response button once, and a negative response by pressing twice. In what follows, we sometimes refer to trials that are accompanied by spatial instructions as the “refined spatial” task, and to trials accompanied by object instructions as the “refined object” task. The modifier “refined” is dropped whenever the context makes it clear which task is under discussion.

Again control tasks were required to minimize the memory component and to subtract from the memory tasks. The controls were identical to the memory tasks except that the retention interval was reduced from 3000 to 250 msec (and the initial fixation cross now remained in view for an additional 2750 msec, so that the total duration of a trial remained constant—see bottom of Fig. 4). In essence, rather than trying to substitute one stage for another—i.e., substitution of a perceptual-decision stage for a memorial-decision stage in Experiment 1—now we simply altered the duration of one stage, the storage stage. This switch from the assumption of stage insertion to one based on stage duration was introduced into chronometric studies of memory 25 years ago by Sternberg (1969). The rationale for the innovation is this: It is difficult to delete an entire processing stage as (as we attempted to do in the control condition of Experiment 1) without affecting other stages in task performance as well, whereas changing just the duration of a stage is more likely to leave the rest of the processing in a task unaffected. The use of this component of Sternberg’s (1969) logic may prove as informative in brain-imaging studies as in strictly behavioral studies.

The sequence of trial events given in Figure 4 depicts the events in all four tasks or conditions of interest—refined spatial memory, refined spatial control, refined object memory, and refined object control. The main differences between the tasks are the instructions (attending to spatial vs. object information) and the duration of the retention interval (3000 vs. 250 msec).

The task alterations of Experiment 2 solve the four general problems that surfaced in Experiment 1: (1) because the objects are so irregular, the likelihood of subjects naming them should be reduced; (2) because the stimuli are identical in the refined spatial and object tasks, differences in activation patterns between the two tasks cannot be tied to the specific stimuli used; (3) because the refined control tasks require *some* memory,

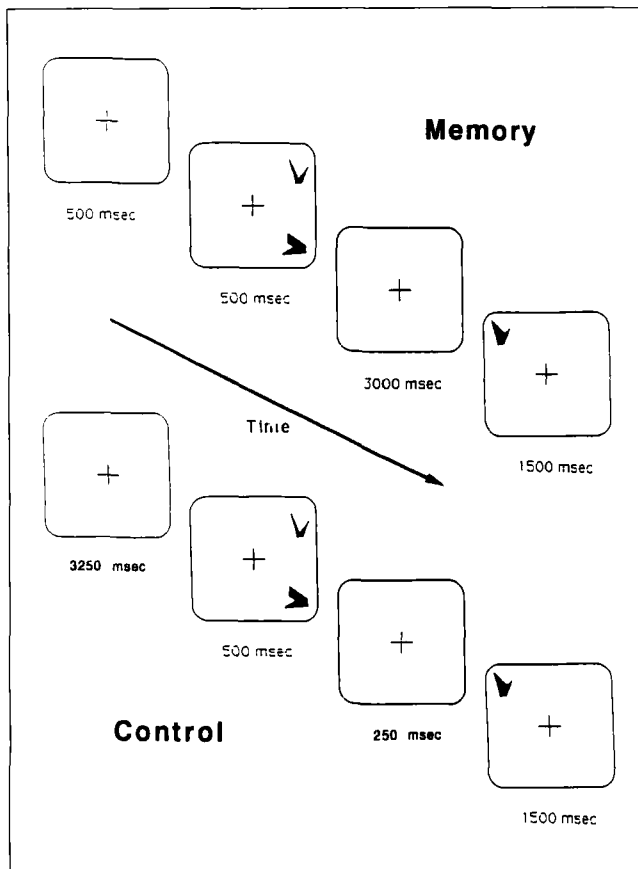


Figure 4. Schematic representation of a trial in the spatial and object tasks in Experiments 2 and 3.

subjects must encode the information, spatial or object, into working memory; and (4) because the visual events in the memory and control tasks are identical, there is no possibility of differential encoding.

A final procedural modification in this experiment is that the same subjects performed in all four tasks. Eight-

een subjects were given 2 blocks of trials on each of the 4 tasks, with the order of conditions counterbalanced across subjects. Again all subjects were right-handed adults who were given practice in all four conditions before being tested in the experiment proper.

Results and Discussion

Behavioral Results

The behavioral findings were similar to those in Experiment 1. Accuracy in the control tasks exceeded that in the memory tasks: for the refined spatial tasks, the difference was 88 vs. 78%; for the refined object tasks, 84 vs. 77% [averaging over spatial and object conditions, $t(17) = 48.70, p < 0.001$]. As in Experiment 1, we also found a difference in the spatial-memory task between performance on distractor trials as a function of whether the probe was relatively far or near from one of the target locations. Subjects correctly rejected 98% of the far distractors, but only 61% of the near ones [$t(17) = 7.85, p < 0.001$].

PET Results

For both refined spatial and object tasks, we created difference images by subtracting activations in the control conditions from those in the respective memory conditions. This time, though, the difference images were analyzed by *region of interest* (ROI) analyses rather than post hoc *t* tests. The ROIs included those regions of activation found significant in Experiment 1, as well as Area 46 because of its significance in studies of working memory. (See the Methods section at the end of the article for details on how these regions were determined.) The top half of Table 2 gives the results for the spatial conditions and the bottom half the results for the object conditions; in each case, the presented coordinates are for the center of mass of the ROI.

Table 2. Stereotactic Coordinates, *t* values, and Corresponding Brodmann Areas for ROIs Showing Significant Activation in Experiment 2 (the Coordinates Are for the Center of Mass in the ROI)

Task	Stereotactic Coordinates			<i>t</i> Statistic	Brain Area
	x	y	z		
Spatial	-32	18	-1	2.96	Right 47 (inferior dorsolateral frontal)
	-34	-52	34	3.67	Right 40/19 (occipitoparietal)
	-29	-2	45	2.93	Right 6 (premotor)
	36	-41	34	2.88	Left 40 (posterior parietal)
	-1	14	43	1.81	32 (anterior cingulate)
	-40	36	22	1.68	Right 46 (mid dorsolateral frontal)
Object	36	-41	34	2.11	Left 40 (posterior parietal)
	48	-59	-10	1.64	Left 37 (inferior temporal)

Consider first the results for the refined spatial task. The right-hemisphere regions that were significant in the spatial task of Experiment 1 are again significant here—Area 47 in prefrontal cortex, Area 6 in premotor cortex, Area 19 in occipital cortex, and Area 40 in parietal cortex (where the latter two areas formed one continuous region of activation). This is good evidence that these regions are indeed involved in the short-term storage of spatial information. Three additional areas also manifested significant (or near significant) activation. One of these additional regions is Area 40 in left-hemisphere, parietal cortex. Although we have found indications of left-hemisphere activity in the spatial task before (see Note 1), this is the first time that such activation has reached significance (partly due to the greater sensitivity of ROI analyses over post hoc tests—see below). This left-hemisphere activation may reflect a “spillover” of the activation from Area 40 in the right hemisphere, which was the single most active region in this study; alternatively, it may reflect *recruitment* of a homologous left-hemisphere region to assist in a particularly demanding version of what is typically a right-hemisphere task (see Smith, Jonides, & Koeppel, 1996). An additional area of activation was the anterior cingulate. This is compatible with the fact that the refined spatial task was more difficult than the spatial task in Experiment 1 (behavioral accuracy averaged 83% in the refined tasks compared to 92% in the Experiment 1 tasks), and that anterior-cingulate activity may reflect overall processing or attentional demand (Pardo et al., 1990). Another additional area of activation is Area 46 in right, prefrontal cortex [the activation here is on the verge of conventional significance by a one-tailed t test, $t(17) = 1.68, p = 0.056$]. Recall that this is the region that is argued to be homologous to the critical area implicated in single-cell studies of spatial working memory (e.g., Wilson et al., 1993). Thus the present results provide something of a bridge to the single-cell studies that used comparable tasks.

Consider now the results for the refined object task. There are only two significant regions of activation, both in the left hemisphere. Hence once more we documented a double dissociation between spatial and object conditions. The two significant regions replicate two found in Experiment 1: Area 40 in parietal cortex and Area 37 in inferotemporal cortex [the latter activation is only of borderline significance— $t(17) = 1.64, p = 0.06$]. The two regions significant in Experiment 1 but not here are Area 32, the anterior cingulate, and Area 6 in premotor cortex (adjacent to Broca’s area).³

How do these findings address the four concerns that motivated this study? First, the fact that the refined object conditions occasioned no significant activation in motor or speech areas suggests that our use of irregular polygons succeeded in reducing verbal coding of the targets. The contrasting activation patterns between the spatial and object conditions now presumably reflect only the difference between spatial and object working

memories, and provide strong evidence that these kinds of memories are mediated by different systems. Second, the fact that we obtained these different activation patterns even though the stimuli were identical in the spatial and object tasks implies that the patterns are in no way artifacts of the specific stimuli used. Third, the fact that we obtained these patterns of activation even when the control or subtraction condition involved some memory encoding and retrieval strengthens our claim that the activation patterns reflect storage processes per se. Fourth, the activation patterns seen in each of the memory-control subtractions cannot be attributed to processing differences for different visual events because the events were identical in memory and control conditions.

Some caveats are in order, though. Because this experiment used ROI analyses whereas the previous one employed less sensitive post hoc tests, we were able to document statistically significant effects for lower activations. With regard to the spatial-memory condition, the regions that were significant in both studies showed an average increase of 3.4% over the control condition in the first experiment, but only 2.5% in the second experiment; likewise, for the object memory conditions, the two regions significant in both studies show a 3.2% effect in the previous study but only a 1.7% effect in the current one.⁴ Although these differences in effect sizes are too small to reach significance, they are suggestive. Their most plausible interpretation is that some of the activation observed in Experiment 1 reflects processes that encode information into working memory, processes that were mitigated in the subtractions of this experiment.

We have emphasized the importance of subtracting out encoding processes from the memory condition so as to get a clearer picture of activation due to storage processes, but we do not mean to imply that encoding and storage necessarily involve different neural structures. For one thing, any active rehearsal processes that occur during the storage interval might amount to a reencoding of the input; for another, passive storage processes might reduce to continued activation of the structures that mediate encoding. In either situation, storage and encoding could involve identical neural structures, in which case subtracting out encoding processes would merely decrease the activation levels in the common structures. There is a bit of evidence for this view that comes from analysis of the activation patterns of just the control conditions (conditions that should reflect encoding but not storage processes). When the object control is subtracted from the spatial control, there is substantial activation in the posterior parietal region of the right hemisphere ($t = 3.9$); this of course is one of the regions that was active in the spatial-memory condition after the spatial control had been subtracted out. Hence this region is active during encoding, and even more active when storage is added to encod-

ing. Similarly, when the spatial control is subtracted from the object control, there is substantial activation in the inferotemporal region of the left hemisphere ($t = 4.9$); this region was active in the object memory condition after the memory control had been subtracted out. So again we have a region that is active during encoding and even more active when storage is added to encoding. Encoding and storage, then, are somewhat intertwined in working memory.⁵

EXPERIMENT 3

Rationale and Tasks

The preceding experiment was premised on the assumption that subjects could attend to spatial information and ignore object information in the spatial conditions, and attend to object information and ignore spatial information in the object conditions. One goal of the present study was to test this assumption directly. The best way to do this is with a strictly behavioral study. A second goal of the present study was to produce a strictly behavioral double dissociation between the spatial and object working-memory tasks.

The tasks used are the refined spatial and object memory tasks—fixation cross, target objects, 3-sec retention interval, and probe for either a spatial or object decision. (No control tasks were needed given the logic of the study.) There are two critical variables, both of which concern only distractor probes, i.e., probes that do not match either target. One variable is the similarity in location of a distractor to the targets (near vs. far); the other factor is the similarity in shape of a distractor to the targets (similar or dissimilar). Both factors were varied in both the spatial and object tasks. The predictions of interest are as follows:

In the refined spatial task, performance should be better on far than near distractors, but be unaffected by the shape similarity of the distractor to the targets.

In the refined object task, performance should be better on distractors that are dissimilar in shape to the targets than on distractors that are similar in shape, but be unaffected by the locational similarity of the distractor to the targets.

Confirmation of these predictions would support the assumption that subjects can selectively attend to either spatial or object information in memory tasks. In addition, it would provide a behavioral double dissociation between our spatial and object working-memory tasks.

As an aid to understanding, Figure 5 presents three sample trials. In the trial schematized on top, the probe is a distractor that is near in location to one of the targets but dissimilar in shape to both of them. In the trial schematized in the middle of the figure, the probe is a distractor that is similar in shape to one of the targets,

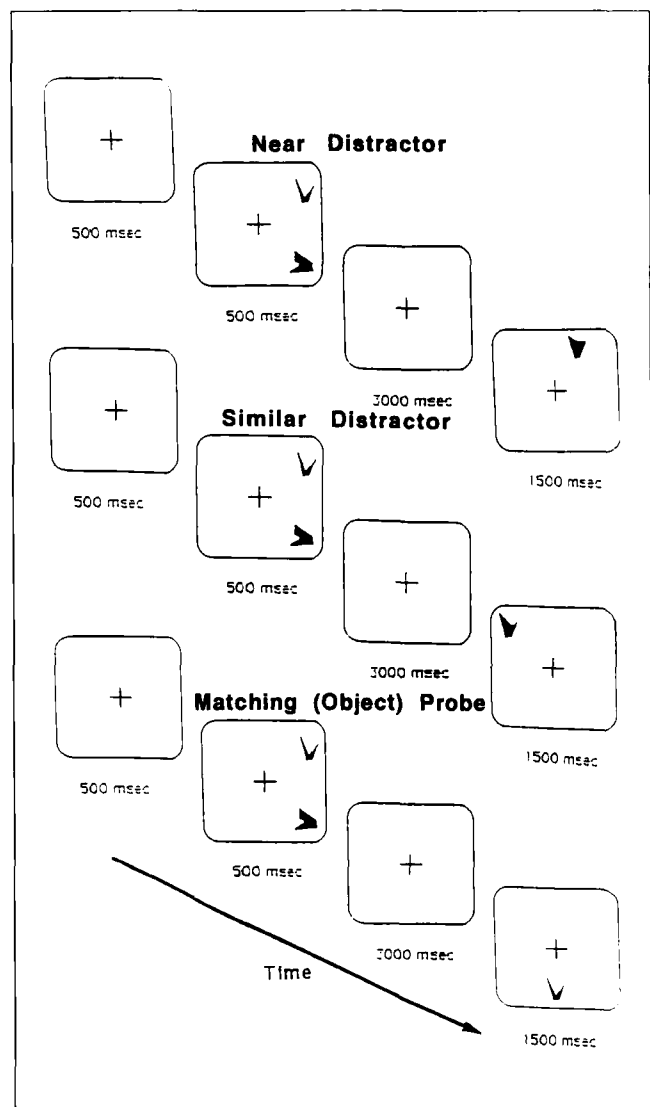


Figure 5. Schematic representation of three sample trials in Experiment 3. See text for explanation.

but dissimilar in location to both of them. According to the predictions of interest, the top trial sequence should lead to relatively poor performance in the spatial task but not in the object task, whereas the middle trial sequence should lead to the opposite pattern of results. The trial at the bottom of Figure 5 illustrates a case in which the probe matches one of the targets in identity. Note that the probe does not match that target (or the other one) in location. This characteristic was true of matching, or affirmative, trials, as we did not want to call the subjects' attention to the information they were trying to ignore.

Figure 5 also makes clear one other methodological point that is needed to appreciate the results (a fuller description is given in the Methods section at the end of the article). Observe that the two targets are relatively far from one another and relatively dissimilar. As a con-

sequence, it was possible to generate two distractor probes that varied in their locational or shape similarity to one of the targets without having the probe-distractor be at all similar (in location or shape) to the other target.

As in the previous experiments, the spatial and object tasks were presented in separate blocks of trials. A total of 10 subjects were given 6 blocks of trials on each of the 2 tasks, with the order of tasks counterbalanced across subjects. The two factors—locational and shape similarity—were manipulated within each block. All subjects were college-age adults and were given practice in both tasks prior to being tested in the experiment proper.

Results and Discussion

Performance was measured in terms of accuracy and latencies. Both sets of data support the predictions of interest, and provide still another double dissociation between the spatial and object tasks.

Figure 6 presents the accuracy results for distractors, with the results for the spatial task graphed on top, and those for the object task graphed on the bottom. Starting with the spatial task, the two bars on the left depict the accuracy for distractors as a function of whether they were relatively near or far to one of the targets. The two bars on the right depict the percent correct for these *same* distractors, now partitioned on the basis of whether they were relatively similar or dissimilar in shape to one of the targets. There is a large effect of the variation in spatial proximity: accuracy on far distractors averaged 98% compared to 63% on near distractors [$t(9) = 5.51, p < 0.001$]. In contrast, these same data show no significant effect of shape similarity: accuracy on dissimilar shape distractors was 80%, on similar shape distractors, 77% [$t(9) = 0.79, p = 0.45$]. This is half the double dissociation. The other half is found in the results for the object task. The bars represent the same entities as before. Now, however, there is little effect of the variation in spatial proximity: accuracy on far distractors is 72 versus 70% on near distractors [$t(9) = 0.92, p = 0.38$]. When these same distractors are partitioned on the basis of shape similarity, however, there is a significant effect: accuracy is 88% on dissimilar distractors versus 58% on similar ones [$t(9) = 6.65, p < 0.001$].

Figure 7 presents comparable data for the reaction times (RTs) of correct responses. Again the data for the spatial task are at the top of the figure whereas those for the object task are at the bottom, and the left pair of bars depict data as a function of spatial proximity whereas the right pair depict the same data partitioned on the basis of shape similarity. Consider first the spatial data. Far distractors were responded to significantly faster than near distractors, 575 vs. 627 msec, respectively [$t(9) = 4.39, p = 0.01$], whereas there was no significant difference when these same distractors were partitioned into dissimilar and similar, 592 vs. 588

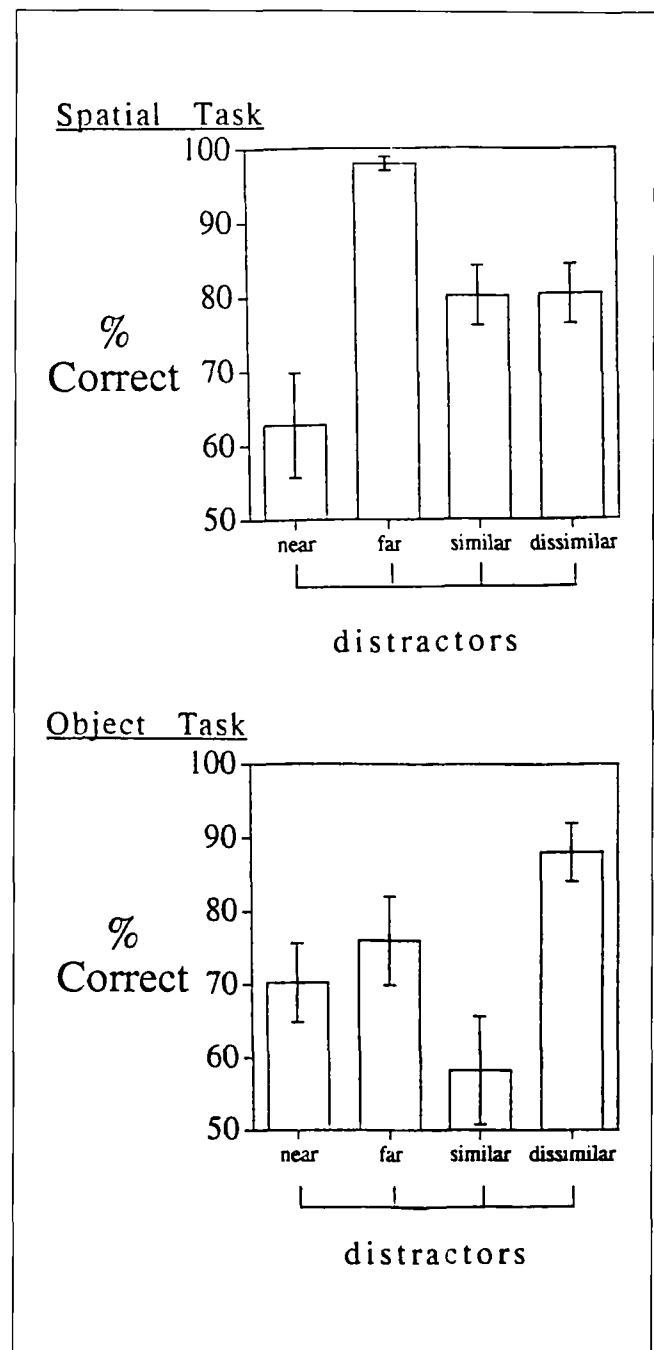


Figure 6. Percent correct response for different kinds of distractors in Experiment 3, separately for spatial task (top) and object task (bottom). Lines through tops of bars indicate standard errors.

msec, respectively [$t(9) < 1$]. The data for the object task reveal the opposite pattern of results. Now there is no effect of whether the distractors are far or near, 665 vs. 656 msec, respectively [$t(9) < 1$], but when the same data are partitioned on the basis of shape similarity, dissimilar distractors yield responses that are significantly faster than similar ones, 649 vs. 702 msec, respectively [$t(9) = 3.22, p = 0.01$]. The RTs therefore, manifest the same kind of double dissociation as do the accuracy data.⁶

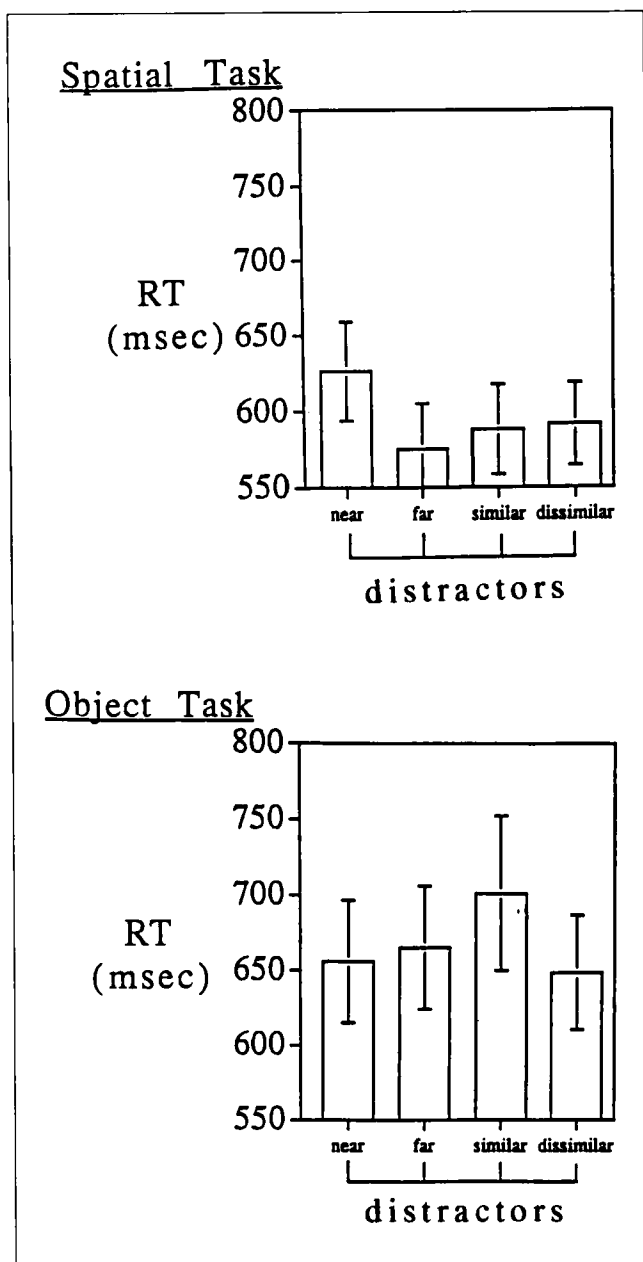


Figure 7. Mean reaction time (in msec) for different kinds of distractors in Experiment 3, separately for spatial task (top) and object task (bottom). Lines through tops of bars indicate standard errors.

It is worth noting that a similar double dissociation has recently been reported by Tresch, Sinnamon, and Seamon (1993). They too used spatial and object working-memory tasks. The spatial task required subjects to remember the location of a single dot during a 10 sec retention interval, whereas the object task required subjects to remember a single object (a simple geometric form) for a 10 sec interval. Rather than manipulate factors intrinsic to the task, Tresch et al. had their subjects perform a secondary task during the retention interval, and varied whether that task involved a movement or a color discrimination. (The rationale for this manipulation was based on single-cell findings indicating that, in visual

perception, location and movement discrimination are mediated by the magnocellular system, whereas shape and color discrimination are mediated by the parvocellular system.) Performance on the spatial working-memory task was disrupted by the movement-discrimination secondary task but not by the color-discrimination task, whereas performance on the object working-memory task was disrupted by color discrimination but not by movement discrimination. This double dissociation dovetails nicely with the one that we just presented. In addition, the results of Tresch et al. show not only that there is a qualitative difference between spatial and object information, but also that this difference in information types holds during the retention interval, not just when the memory decision is made.⁷

GENERAL DISCUSSION

In what follows, first we briefly summarize our major findings, then consider the relation between our results and other relevant findings, and finally outline a model of spatial and object working memory.

Summary

The current results provide strong evidence for the claim that spatial and object working memories are mediated by qualitatively different systems. Three kinds of results are important to this claim.

First, our PET measures showed a double dissociation between the brain regions active during the storage stage of a spatial working-memory task and those active during the storage stage of an object working memory task; the spatial-task regions are localized almost exclusively in the right hemisphere, whereas the object-task regions are left-hemisphere areas. Moreover, we obtained this kind of double dissociation in two separate studies (Experiments 1 and 2), which used different stimuli, tasks, and control conditions.

Second, the known functionality of the specific regions activated provides further support for the claim of separate systems. The spatial task, but not the object one, led to activation in Area 40 of right parietal cortex, a region known to be involved in the formation of spatial representations. In contrast, the object task, but not the spatial one, was associated with activation in the infero-temporal cortex, a region known to be involved in the formation of shape representations. The results about specific regions of activation proved more robust for spatial than object working memory in that every site of spatial-task activation found significant in Experiment 1 was also significant in Experiment 2, whereas only two of four sites of object-task activation in Experiment 1 were also significant in Experiment 2 (and one of these was only marginally significant in Experiment 2). However, the two Experiment 1 regions in the object task that did not approach significance in Experiment 2 were

those whose known functionality has little to do with object or shape memory (the anterior cingulate and part of premotor cortex). All things considered, the findings about specific regions leave little doubt that fundamentally different kinds of information comprise the contents of spatial and object working memory. Any attempt to account for working memory in terms of a uniform (e.g., propositional) representation—the kind of account suggested by Pylyshyn's (1981) analysis of imagery, for example—seems doomed from the start.

Third, the results from our strictly behavioral study, Experiment 3, converged with the findings from our PET studies. This convergence is important because the behavioral and PET studies have complementary strengths and weaknesses. For example, although only the PET experiments provide evidence about underlying neural bases of memory, only the behavioral experiment permits a subtle within-block variation in type of distractor.⁸

Relation to Other Findings

Other Imaging Studies

Although no other experiments have used imaging techniques to study the contrast between spatial and object working memory, a couple of previous imaging studies have looked at one of these kinds of memories. McCarthy, Blamire, Puce, Nobre, Bloch, Hyder, Goldman-Rakic, and Shulman (1994) used functional magnetic resonance imaging to examine spatial working memory. They required their subjects to determine whether or not a currently presented visual form was in the same spatial position as that of a previously presented form. McCarthy et al. (1994) found heightened activation in prefrontal cortex, Area 46. This, of course, is one of the regions that we found to be active in spatial working memory. The other relevant study involves only object working memory (Petrides, Alivisatos, Evans, & Meyer, 1993). Subjects were presented a series of cards, each of which contained the same eight abstract designs; they had to point to a different design on each card, until all eight designs had been selected. Accurate performance therefore required that subjects remember their previous selections, and presumably some of this load was maintained in object working memory. Petrides et al. found numerous cortical regions to be active in this task, even after subtracting out activation from a control condition. One such region was Area 46 (particularly in the right hemisphere). This does not fit with the findings for our object task. However, other regions of activation reported by Petrides et al. (1993) do correspond to some of our findings for the object task, including activation in the left parietal cortex and the anterior cingulate.

Although the commonalities with our results are encouraging, we do not want to make too much of these comparisons. Looking for common activation sites across studies is somewhat problematic because the tasks employed in these studies differ in important ways. In both

our spatial- and object-memory tasks, a relatively small amount of information (two or three items) was presented briefly (less than half a second), and was followed by a blank retention interval during which subjects may have actively engaged in some form of rehearsal process. In contrast, in the tasks of McCarthy et al. (1994) and Petrides et al. (1993), a relatively large amount of information (at least eight items) was presented for a matter of seconds, and the retention interval was filled with the presentation of other items. In these latter studies, it is quite possible that (1) more than working memory was required (because of the larger memory load), (2) the positions or objects were verbally represented (because of the lengthy encoding intervals), and (3) active maintenance processes were mitigated (because of the filled retention intervals). The impact of these differences between tasks needs to be sorted out by future research.

Single-Cell Studies

At the outset we described a single-cell study by Wilson et al. (1993) that showed that a spatial working-memory task in monkeys led to activation in the principal sulcus region of prefrontal cortex, whereas an object working-memory task resulted in activation in a different prefrontal region. At a general level, our results agree with those of Wilson et al. (1993) in showing a double dissociation between spatial and object working memory. However, only in our case is the dissociation lateralized—spatial in the right, object in the left. Though this is a very salient difference, it can be assimilated to the well-known tendency for the human brain to show more lateralization of function than do the brains of other primates (e.g., Kolb & Whishaw, 1990).

At a more specific level, we may compare the two studies with respect to the particular areas activated in prefrontal cortex (the only region that was examined by Wilson et al., 1993). A point of commonality appears to be Area 46. It was activated in our refined spatial task, and it appears to be homologous to the principal sulcus region in monkey brain that was the locus of active neurons in the spatial task of Wilson et al. (Goldman-Rakic, 1993, personal communication). But there are also two important differences between the prefrontal regions implicated in our study and in that of Wilson et al. (1993). First, Wilson et al. (1993) found that neurons in a region inferior to the principal sulcus were active during their object-memory task, whereas we did not find any prefrontal area to be active during our object-memory task (a premotor area was active in Experiment 1 but not in Experiment 2). It is not clear what to make of this difference, particularly since our results for the object-memory task were not that robust. Second, we found that Area 47 (right hemisphere) was activated in the spatial-memory task, whereas there is no indication of neural activation in a homologous region of monkey brain in the Wilson et al. (1993) study. As suggested

earlier, this area may be associated with some sort of visual-spatial rehearsal, a strategy that may be more available to humans than monkeys.

Outline of a Cognitive Model of Spatial and Object Working Memory

Earlier, we suggested two ways—passive storage and active rehearsal—in which the areas activated in our spatial and object tasks might be configured into circuits. Regardless of exactly how these circuits are configured, they must at some level contain networks of neurons that encode and store spatial or object information. These are the crucial computations that mediate performance in spatial and object working-memory tasks, and any theoretical description of working memory must include a detailed account of them. We know of no data that indicate how these computations are accomplished in humans, but the details of the single-cell findings of Goldman-Rakic and her colleagues (e.g., Funahashi, Bruce, & Goldman-Rakic, 1989) suggest a possibility. In what follows, we use these findings to outline a *cognitive* (computational) model of how information is encoded and stored in spatial and object working memory.

Two findings of Goldman-Rakic and her colleagues are particularly important for our concerns:

1. There are sets of neurons (in the principal sulcus region) that are active during the retention interval in a spatial task. Some of these neurons are also active when the target information is initially presented, whereas others are active only during the delay period.
2. Some of these neurons are location-specific. That is, for some specific locations in the visual field the neurons of interest have a high firing rate, whereas for other locations their activity does not exceed baseline. These neurons are broadly tuned, however. For example, a neuron that produces its maximum firing rate to a target presented at 135° will also produce above-baseline activity to a target presented at 90° or 180°.

To use these findings to constrain a model of human spatial working memory, we assume that the short-term storage of location information is mediated by the activation of a set of *location-specific memory units* that are broadly tuned. Thus, when an object is presented at 135°, those location-memory units that are maximally tuned for that spatial position will be activated to their maximal extent, whereas units that are tuned to neighboring positions will be moderately activated. The result is that each possible location in a subject's visual field is associated with a distinctive pattern of activation across a set of location-memory units. As a consequence of this population-coding scheme, locations that are close to one another in a subject's visual field will be represented by similar patterns of activation of location-memory units. This provides a natural account of the main behavioral finding in our spatial tasks, namely that subjects were

more accurate in rejecting far distractors than near ones (Experiments 1–3). The activation pattern of a near distractor would be similar to that of one of the targets, and presumably the probability of correctly rejecting a distractor increases with its absolute dissimilarity from the targets. This coding scheme also makes sense of a second behavioral finding: the effect of spatial proximity arises only when the subject is attending to spatial information (Experiment 3), for only then is the subject activating location-memory units.

Consider now how this model can handle the effects of the duration of the retention interval (recall that in Experiment 2, there was a variation in retention interval from 250 msec to 3 sec when comparing the control to the memory condition). The most obvious assumption to add to the model would be that the activation of location-memory units decays with time. But this assumption seems at odds with the single-cell results of Funahashi et al. (1989), who, surprisingly, found no decline in neuronal activation in a spatial-memory task as the retention interval increased. The intervals studied by Funahashi et al., though, were 3 and 6 sec, whereas ours were 250 msec and 3 sec. Our interval of 250 msec is so brief that perhaps location-encoding units were still active throughout the interval (along with location-memory units); if so, in effect there would have been decay in activation as the retention interval lengthened from 250 msec to 3 sec (because the location-encoding units would not be active throughout the 3-sec interval). The impact of such decay would be that the activation pattern for a stored target no longer would be identical to that of a matching probe. Consequently, subjects would have to make their decision criteria for a match more lenient, allowing positive responses even when the probe was somewhat dissimilar to a target. The upshot of this criterion change would be that many near distractors would now pass criterion, leading to incorrect positive responses. In sum, a decline in activation, coupled with a shift in decision criterion, leads to the prediction that accuracy should decline with the duration of the retention interval, with this decline being greater for near than far distractors. This is exactly what we found in Experiment 2.⁹

A similar story applies in modeling object working memory. Now the short-term storage of object information would be mediated by the activation of a set of *shape-specific memory units*, where each unit is tuned to a part of an object, or a feature of the entire object. Two objects that are similar in shape will be represented by similar sets of shape-memory units. This explains why, in Experiment 3, subjects were more accurate in rejecting dissimilar distractors than similar ones, and why this finding arises only when the subject is attending to object information (i.e., only then is the subject activating shape-memory units). Again we can accommodate the effects of retention-interval duration demonstrated

in Experiment 2 by assuming that the activation of the memory units decays with time. After a few seconds delay, the activation pattern for a stored target no longer would be identical to that of a matching probe, leading subjects to make their decision criteria more lenient, which in turn leads to a decrease in accuracy, particularly on similar distractors.

The preceding models account for all of our behavioral findings, but this is not much of a theoretical feat given the small number of findings and the richness of our assumptions. However, there is another virtue of the models—their components can readily be reduced to plausible neuronal events. More generally, if we desire that a cognitive model not only explain behavioral data but also map onto brain events, then we will be best off using whatever relevant knowledge we have of brain function to constrain the assumptions of the model in the first place (for a related argument, see Kosslyn & Intriligator, 1992).

EXPERIMENT 1 METHOD

Subjects

The two groups of subjects were drawn from the University of Michigan community. All subjects who volunteered to participate for pay were right-handed and of college student age. The group that performed the spatial working-memory task consisted of 11 males and 7 females; the group that performed the object working-memory task consisted of 4 males and 8 females. None of the subjects was taking medication, and all reported good health. No subject was tested in more than one experiment reported in this article.

Tasks and Materials

Spatial Tasks

As described in part earlier, a trial in the spatial-memory task included the following four events:

1. a fixation cross presented for 500 msec;
2. three black dots, each with a diameter of 0.7 cm, presented for 200 msec on the circumference of an imaginary circle; the imaginary circle was centered on the fixation cross and had a diameter of 14°;
3. the fixation cross alone presented for 3000 msec; and
4. an unfilled probe circle, 1.4 cm in diameter, presented for 1500 msec, which, with probability 0.5, encircled the location of one of the three target dots.

Subjects indicated whether or not they thought the probe encircled a previous dot location by pressing with their right hand a computer mouse, either once or twice. For the first 10 subjects, a double press indicated an affirmative response whereas a single press indicated a negative response; for the last 8 subjects the assignment of button presses to affirmative and negative decisions

was reversed. Subjects had to make their response within 1500 msec in order for it to be recorded. A trial in the spatial control task was identical to the preceding except that the order of trial events (2) and (3) was reversed, and the target dots remained in view while the probe circle was presented. Subjects indicated whether or not the probe encircled a dot location by pressing a mouse with their right hand once or twice; again, for the first 10 subjects a double press indicated an affirmative response, while for the last 8 subjects a single press indicated an affirmative response.

To select the targets for a particular trial, 3 different positions were selected from the 360 possible degrees on the imaginary circle, with the restriction that each pair of dots had to be at least 40° apart. The three dots *appeared* to be a random selection, as attested to by pilot subjects. Also, our experimental subjects seemed to treat the dots independently, as only 1 of the 18 subjects reported paying any attention to the relation between the dots. The probe circle was either centered directly over the previous location of a dot, or it missed the nearest dot location by at least 15°.

For each task, 20 trials were grouped together to form a block. The duration of a block was 1 min 54 sec (5200 msec per trial plus 500 msec per trial as an intertrial interval), and there was a 14 min interval between successive blocks. There were 2 blocks for the spatial-memory condition, 2 for the spatial-control condition, and 2 for a third condition that is largely irrelevant to the present paper. For the last 8 subjects tested, the order of these 3 conditions was counterbalanced across subgroups of 4 subjects.

Object Task

A trial in the object-memory task included the following four events:

1. a fixation cross presented for 500 msec;
2. two novel target objects presented for 500 msec on either side of where the fixation cross had been;
3. just the fixation cross presented for 3000 msec; and
4. a probe object presented for 1500 msec in the position previously occupied by the fixation cross, which, with probability 0.5, matched one of the two target objects.

Subjects indicated whether or not they thought the probe object matched either of the target objects by pressing a mouse once (affirmative) or twice (negative) with their right hand. Again subjects had to make their response within 1500 msec for it to be recorded. A trial in the object-control task was identical to the preceding except that the order of trial events (2) and (3) was reversed, and the target objects remained in view when the probe object was presented.

Each object consisted of an interior shape included within an exterior shape, with some connecting lines

between the two shapes. The shapes and connecting lines were constructed by using options in MacPaint. A total of seven different shapes was used: a square, a circle, a diamond, a "tall" rectangle, a "flat" rectangle, a "tall" ellipse, and a "flat" ellipse. These seven forms were used as both interior and exterior shapes, with their size of course varying with the role they played. The lines connecting the interior to the exterior shape were of eight different kinds: a horizontal line connecting leftmost (or rightmost) points of the interior and exterior shapes, a vertical line connecting the topmost (or bottom-most) points of the interior and exterior shapes, and four diagonal lines that terminated on an edge or end of the interior shape. Since each one of these eight lines could be present or absent (and since there were seven different shapes), the total number of possible objects was $2^8 \times 7 \times 7$. The number actually used was far less because, among other constraints, we used 10 different objects during a block of trials corresponding to a scan. Still, from the subject's perspective there were an unlimited number of objects. To select the objects for a particular trial, two different objects from all possible ones were sampled. Care was taken to ensure that the two objects selected were not too similar to one another, which could have led to some kind of abstract coding (e.g., "objects consisting of only rectangular shapes"). When the probe did not match either of the targets, it was drawn at random from the remaining set of objects. No attempt was made to control the similarity of distractor probes to the target objects.

Again, for each task, 20 trials were grouped together to form a block. The duration of a block was 1 min, 55 sec (5500 msec per trial, plus 250 msec per trial as an intertrial interval), and there was a 14 min interval between successive blocks. There were three blocks for the object-memory condition, three for the object-control condition, with the order of the two conditions being counterbalanced over subgroups of six subjects. Before participating in the experiment proper, subjects were given three practice blocks of trials on each condition of interest.

Computer Control

All tasks were conducted with a Macintosh IIfx computer with a 14-inch monitor, using a program written in C to present stimuli and record subjects' responses. The computer rested on a gantry approximately 18 inches from the subject's eyes; the screen was tilted down to face the subjects.

PET Procedure

The PET machine was a Siemens/CTI-931/08-12. The camera produced 15 contiguous slices that were 6.75 mm apart (center to center); the reconstructed axial resolution was 10 mm FWHM. The PET machine was in

a suite designed specifically for PET use, and the same background conditions obtained in all testing—the lights were dimmed and there was no conversation or intrusive noise. Subjects first gave informed consent, then were familiarized with the PET apparatus prior to the experiment proper. Each subject had an intravenous catheter inserted into the left arm to receive the injections of radioactive tracer, and after that they participated in the practice blocks. They were then positioned in the scanner with tape applied from the head holder to their foreheads to constrain head movement. The experimental protocol consisted of six scans, each corresponding to a block of trials.

Each scan consisted of 20 trials, with the first 3 trials presented prior to the injection of the radionuclide (the total duration of these 3 trials was approximately 15 sec). Immediately following these trials an intravenous bolus injection of 66 mCi of oxygen-15-labeled water was administered, after which approximately 15 sec elapsed before the radionuclide reached the brain. Trials continued to be administered during the interval. Recording of activity began 5 sec after the count rate was observed to increase above the background level and continued for 60 sec thereafter. Injections for subsequent scans were separated by 14-min intervals, permitting the oxygen-15 to decay to an acceptable background level.

PET Data Analysis

The PET images for each subject were transformed to a stereotaxic coordinate system (Minoshima, Berger, Lee, & Mintun, 1992; Minoshima, Koeppe, Mintun, Berger, Taylor, Frey, & Kuhl, 1993), and linearly standardized to an atlas brain (Talairach & Tournoux, 1988). After normalizing pixel values for global flow rate differences among scans (Fox, Fox, Raichle, & Burde, 1985), the data were averaged across the subjects in a condition, giving mean and variance values for each condition. The average image for each control condition was subtracted from that of its corresponding memory condition to reveal differences in activation between these conditions. The difference images were then analyzed for statistical significance on a voxel-by-voxel basis, using *t*-statistics, followed by a multiple-comparison adjustment based on the Bonferroni method (Friston, Frith, Liddle, & Frackowiak, 1991; Worsley, Evans, Marrett, & Neelin, 1992). A one-tailed adjusted value of $p < 0.05$ was used as a criterion for reliability (unless otherwise noted).

EXPERIMENT 2 METHOD

Subjects

A group of 18 subjects, drawn from the University of Michigan community, volunteered to participate for pay. There were 8 males and 10 females, all right-handed and of student age. None of the subjects was taking medication, and all reported good health.

Tasks and Materials

As described earlier, a trial in either the spatial-memory or object-memory task included the following four events:

1. a fixation cross presented for 500 msec;
2. two target objects, each an irregular polyhedron, presented for 500 msec in random positions around an imaginary circle;
3. the fixation cross presented for 3000 msec; and
4. a probe object presented for 1500 msec.

In the spatial-memory conditions, subjects were instructed to decide whether or not the probe object was in the same position as one of the target objects. In the object-memory conditions, subjects were instructed to decide whether or not the probe object was identical to either of the target objects. In all conditions, subjects were instructed to respond quickly while making accuracy their first priority. Also, in all conditions, subjects responded affirmatively by pressing a mouse once, and negatively by pressing the mouse twice (both responses being made with the right hand). A trial in the spatial- or object-control task was identical to that in the corresponding memory task, except that the duration of the third trial event—the retention interval—was reduced from 3000 to 250 msec, and the duration of the initial trial event—the fixation period—was increased from 500 to 3250 msec.

All subjects served in all four conditions. This made it necessary for subjects to selectively attend only to spatial information in the spatial tasks, and only to object information in the object tasks. An effort was made not to call subjects' attention to the irrelevant dimension. Specifically, in the spatial tasks a probe never matched a target in identity, and in the object tasks a probe never matched a target in position. One further comment about distractor probes: In the spatial task, half the distractor probes were relatively near to one of the target objects and half were relatively far; near probes were between 15° and 25° from the closer target, whereas far probes were between 40° and 50° from the closer target. There was no corresponding variation for distractors in the object task.

The objects used were "Attneave" shapes. Attneave generated these shapes by (1) randomly selecting n points from a 100-by-100 matrix of points, and (2) connecting the n points to form a shape (see Attneave & Arnoult, 1956). We used only 4-point and 6-point shapes that had previously been employed by Vanderplas and Garvin (1959). The objects were divided into 40 sets of 4 figures each. The objects within a set had the following similarity relations (illustrated by the sample set in Fig. 8). One object, call it "A," served as a kind of starting point. The remaining three objects—"B," "C," and "D,"—were selected such that B was more similar to A than C was, and the similarity between A and C exceeded that

between either C and D or B and D. A and D served as targets, whereas B and C served as similar and dissimilar distractors, respectively, where similarity is defined with respect to A. In this way, we ensured that when a dissimilar distractor was used (a C object), it was not similar to the other target.

To document these relations within each set, we obtained similarity ratings from an independent group of 15 subjects (all University of Michigan undergraduates). Pairs of objects from within a potential set of 4 were presented, and subjects rated their similarity on a 7-point scale. The objects were presented at the same rate and size as during the actual experiment. After averaging the similarity ratings over subjects, we checked the similarity relations within a set, and kept only those sets that had the desired relations and where each relevant similarity difference was at least one scale unit in magnitude.

For each of the four tasks—spatial memory, spatial control, object memory, and object control—20 trials were grouped together to form a block. The duration of a block was 1 min 55 sec (5500 msec per trial plus 250 msec per trial as an intertrial interval). Again there was a 14-min interval between successive blocks. There were two blocks for each of the four conditions, with the order of the conditions counterbalanced over subgroups of four subjects.

Computer Control

The computer control of this study was the same as in the previous experiment.

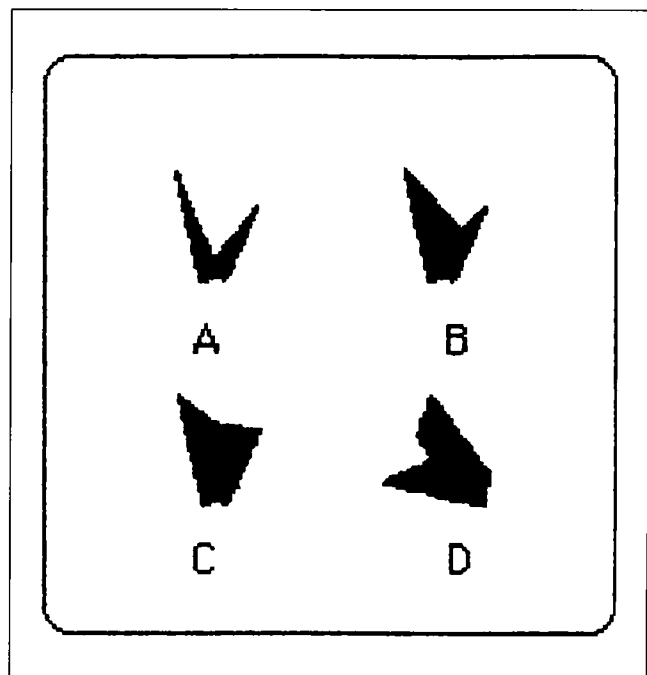


Figure 8. One set of four objects used in Experiments 2 and 3. See text for explanation.

PET Procedure

A new PET machine was used, a Siemens ECAT EXACT-47. The camera produced contiguous slices that were 3.375 mm apart (center to center); the reconstructed axial resolution was 10 mm FWHM. Other aspects concerning the use of PET were the same as in Experiment 1, except that 50 mCi (rather than 66 mCi) of oxygen-15-labeled water was administered for each scan, since there were now 8 scans per subject rather than 6.

PET Data Analysis

Difference images were created in the same manner as in Experiment 1, but the statistical analysis of the images differed. Rather than using post hoc *t* tests, we employed the results of Experiment 1 to define regions of interest (ROIs). Some ROIs were based on the results of the spatial task in Experiment 1, other ROIs were based on the results of the object task in Experiment 1, and ROIs for Area 46 were determined by use of the Talairach and Tournoux (1988) Atlas.

A spatial-based ROI was determined as follows:

1. First, for each of the four areas that produced significant activation in Experiment 1, we selected the voxel associated with the peak or maximum activation; this voxel served as a "seed" for the ROI; and
2. We allowed the ROI to "grow" by adding to the seed voxel any adjacent voxel whose activation had a *z* value of 4.0 or greater. (Each of the ROIs had at least 11-12 voxels, which means each ROI was at least 0.125 cm³.)

A total of four distinct spatial-based ROIs were produced, corresponding to the four significant areas of activation. In like manner, four object-based ROIs were determined. The only changes from the above procedure were that the seed voxels were, of course, based on the object task of Experiment 1, and that the adjacent voxels added onto the seed had to be associated with a *z* value of only 3 or greater (due to the fact that the obtained significance levels were lower in the object than the spatial task in Experiment 1). To create ROIs for Area 46 (for both hemispheres), we proceeded as follows:

1. First, on the basis of the Talairach and Tournoux (1988) Atlas, we located Area 46, and then divided it into superior and inferior halves. We did this division because results that we have obtained in other PET experiments (not reported here) suggest that the superior portion of Area 46 is more active in spatial tasks than is the inferior portion.
2. For each section, we calculated the voxel corresponding to the area's center of mass.
3. We then established a spherical region centered on the center-of-mass voxel. This sphere had a diameter of 11.25 mm (roughly 5 voxels); use of a larger diameter would have included voxels that were outside of Area 46, whereas use of a smaller diameter would have ex-

cluded numerous voxels that were in Area 46. The resulting region contained about 65 voxels, and served as the ROI.

The resulting 10 ROIs (4 spatial-based, 4 object-based, and 2 Atlas-based) were applied to the activation results of both the refined spatial and object tasks.

Finally, we performed *t* tests on all ROI's. In each test, we determined the average difference in activation between the memory and control tasks, computed over all voxels in the ROI. As expected, in the spatial task, activation in the upper portion of Area 46 was closer to significance than activation in the lower portion of Area 46. However, the difference in *t* is quite small. The entry for the Area 46 ROI in Table 2 is for the upper portion of 46, and has a significance level of $p = 0.056$; if the ROI is instead based on all of Area 46 the level of significance is reduced only marginally to $p = 0.068$.

EXPERIMENT 3 METHOD

Subjects

The subjects were 10 University of Michigan students, who volunteered to participate for pay. All were right-handed.

Tasks and Materials

The two tasks were the spatial-memory and object-memory tasks developed in Experiment 2 (there was no need for control tasks given the logic of the study). One important change from the Experiment 2 methodology concerned the construction of trial blocks. Now 40 trials were grouped together to form a block, and 6 blocks of each task were presented to each subject. Another change from Experiment 2 concerned how subjects were instructed about the relative importance of the speed and accuracy of their responses. Subjects were now told "please answer as quickly as possible, but remember that accuracy is very important." They were further informed that they would receive a bonus for each correct response made in less than a second, but only half a bonus for a correct response that took longer than a second.

Computer Control

The experiment was conducted with a Macintosh IICI computer with a 14-inch screen, using a program written in C to present stimuli and record subjects' responses. The subject sat approximately 18 inches from the screen.

Acknowledgments

The research reported in this paper was supported by grants from the Office of Naval Research, the James S. McDonnell Foundation, and the Department of Energy (DE-FG02-87ER60561). Part of the research was performed while both

Smith and Jonides were supported by training grants from the James S. McDonnell Foundation.

Reprint requests should be sent to Dr. Edward E. Smith, Department of Psychology, University of Michigan, 525 East University, Ann Arbor, MI 48109-1109.

Notes

1. While the four areas of activation just discussed are in the right hemisphere, it is worth noting that there was also increased activation in the homologous areas in the left hemisphere. However, none of the left-hemisphere activation levels reached significance. In addition, the correlations between the activation levels in left-hemisphere regions and accuracy in the behavioral task were essentially 0 (average $r = 0.01$), whereas the comparable correlations for the right-hemisphere regions were substantially larger (average $r = 0.33$). All things considered, performance in the spatial memory task seemed to be mediated more by right-hemisphere than left-hemisphere activation.

2. The results described in the text all concern activations. There were also some significant *deactivations*, i.e., reliable decreases in activation in the memory condition compared to the control condition. In the spatial conditions, two sites of deactivation reached significance (as manifested by a z of 4.1 or better), both of which appeared to be in the posterior cingulate. (The x , y , and z coordinates for the two sites are 1, -51, 20 and -1, -26, 36.) In the object conditions, only one deactivation reached significance: again the activation appeared to be in the posterior cingulate (with coordinates -3, -51, 22). The most straightforward interpretation of these results is that the neural basis of short-term storage involves an inhibition of the posterior cingulate. This interpretation is strengthened by an additional finding in the spatial experiment. In that study, some of the subjects also participated in a rest condition in which they merely attended to a fixation cross. When the activation pattern for this condition is subtracted from that in the memory condition, once more a significant deactivation is found in the posterior cingulate (with coordinates 1, -53, 20).

3. There was again evidence for deactivation of the posterior cingulate in memory conditions (see Note 2). In the spatial conditions, the deactivated site reached significance (with coordinates 3, -51, 27). In the object conditions, the deactivation of the posterior cingulate failed to reach significance, though this site was the most deactivated one found (with coordinates 1, -60, 27).

4. The percentage increases are based on the maximally active voxel in the region.

5. We are indebted to Patricia Reuter-Lorenz for suggesting this analysis of the control conditions.

6. The data described in the text are for distractor probes only. For probes that matched a target, the accuracies were 87 and 70% in the spatial and object tasks, respectively. The RTs for correct responses to such probes were 580 and 661 msec in the spatial and object task, respectively.

7. Our behavioral study, like that of Tresch et al. (1993), used simple spatial and object tasks that seem to involve little else than spatial and object working memory. In contrast, prior behavioral studies of spatial vs. object working memory have used more complex tasks that may have recruited some unwanted kinds of processing. This problem is particularly true of Baddeley and Lieberman (1980), and, to some extent, may apply to an otherwise elegant study by Logie (1986). Still, all of these behavioral studies converge in showing a double dissociation between spatial and object tasks.

8. Another kind of converging study would involve brain-damaged human patients, demonstrating that damage in one cortical region is associated with an impairment only in spatial working memory, whereas damage in another area is associated with an impairment only in object memory. We know of no such results with tasks that tap *only* short-term storage. However, this sort of double dissociation has been obtained in tasks that clearly require long-term storage, and perhaps working memory as well (e.g., generate a visual image of a dog from memory—Levine, Warach, & Farah, 1985).

9. There is an alternative way to model the effects of retention interval. The overall level of an activation pattern may remain constant through the retention interval, but the pattern itself may become progressively distorted as the interval lengthens because the variance associated with each memory unit is increasing as the interval lengthens. One reason why the variance may increase with the interval is that the subject is attempting to "rehearse" the neural pattern associated with the input, and each successive rehearsal may blur the activation pattern a bit.

REFERENCES

- Alivisatos, B., & Milner, B. (1989). Effects of frontal or temporal lobectomy on the use of advance information in a choice reaction time task. *Neuropsychologia*, 27, 495-503.
- Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.
- Attneave, R., & Arnoult, M. D. (1956). Methodological considerations in the quantitative study of shape and pattern in perception. *Psychological Bulletin*, 53, 452-471.
- Baddeley, A. D. (1986). *Working memory*. Oxford: Oxford University Press.
- Baddeley, A. D. (1992). Working memory. *Science*, 255, 566-569.
- Baddeley, A. D., & Lieberman, K. (1980). Spatial working memory. In R. S. Nickerson (Ed.), *Attention and performance VIII* (pp. 521-539). Hillsdale, NJ: Erlbaum.
- Carpenter, P. A., Just, M. A., & Shell, P. (1990). What one intelligence test measures: A theoretical account of the processing in the Ravens Progressive Matrices Test. *Psychological Review*, 97, 404-431.
- Egley, R., Driver, J., & Rafal, R. D. (1994). Shifting visual attention between objects and locations: Evidence from normals and parietal lesion subjects. *Journal of Experimental Psychology: General*, 123, 161-177.
- Farah, M. J. (1991). Patterns of co-occurrence among the associative agnosias: Implications for visual object recognition. *Cognitive Neuropsychology*, 8, 1-19.
- Fox, P. T., Fox, J. M., Raichle, M. E., & Burde, R. M. (1985). The role of cerebral cortex in the generation of saccadic eye movements; a positron emission tomography study. *Journal of Neurophysiology*, 54, 348-368.
- Friston, K. J., Frith, C. D., Liddle, P. F., & Frackowiak, R. S. J. (1991). Comparing functional (PET) images: The assessment of significant change. *Journal of Cerebral Blood Flow & Metabolism*, 11, 690-699.
- Funahashi, S., Bruce, C. J., & Goldman-Rakic, P. S. (1989). Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *Journal of Neurophysiology*, 61, (2), 331-349.
- Goldman-Rakic, P. S. (1987). In F. Plum (Ed.), *Circuitry of primate prefrontal cortex and regulation of behavior by representational memory*. *Handbook of physiology: The nervous system*. Bethesda, MD: American Physiological Society. 373-417.
- Grady, C. L., Haxby, J. V., Horwitz, B., Ungerleider, L. G.,

- Shapiro, M. B., Carson, R. E., Herscovitch, P., Mishkin, M., & Rapoport, S. I. (1991). Dissociation of object and spatial vision in human extrastriate cortex. *Journal of Cognitive Neuroscience*, 4, 23-24.
- Jonides, J. (1995). Working memory and thought. In D. Osherson & E. E. Smith (Eds.), *Invitation to cognitive science, Vol. 3, Thinking* (2nd ed.). Cambridge: MIT Press. In press.
- Jonides, J., Smith, E. E., Koeppe, R. A., Awh, E., Minoshima, S., & Mintun, M. A. (1993). Spatial working memory in humans as revealed by PET. *Nature (London)*, 363, 623-625.
- Just, M. A., & Carpenter, P. A. (1992). A capacity theory of comprehension: Individual differences in working memory. *Psychological Review*, 99, (1), 122-149.
- Kolb, B., & Whishaw, I. Q. (1990). *Fundamentals of human neuropsychology*. New York: Freeman.
- Kosslyn, S. M., Alpert, N. M., Thompson, W. L., Maljkovic, V., Weise, S. B., Chabris, C. F., Hamilton, S. E., Rauch, S. L., & Buonanno, F. S. (1993). Visual mental imagery activates topographically organized visual cortex. *Journal of Cognitive Neuroscience*, 5, 263-287.
- Kosslyn, S. M., & Intriligator, J. R. (1992). Is cognitive neuropsychology plausible? The perils of sitting on a one-legged stool. *Journal of Cognitive Neuroscience*, 4, 196-106.
- Levine, D. N., Warach, J., & Farah, M. J. (1985). Two visual systems in mental imagery: Dissociation of 'what' and 'where' in imagery disorders due to bilateral posterior cerebral lesions. *Neurology*, 35, 1010-1018.
- Logie, R. H. (1986). Visuo-spatial processing in working memory. *Quarterly Journal of Experimental Psychology*, 38A, 229-247.
- McCarthy, G., Blamire, A. M., Puce, A., Nobre, A. C., Bloch, G., Hyder, F., Goldman-Rakic, P., & Shulman, R. G. (1994). Functional MR imaging of human prefrontal cortex activation during a spatial working memory task. *Proceedings of the National Academy of Sciences, USA, Neurobiology*.
- Minoshima, S., Berger, K. L., Lee, K. S., & Mintun, M. A. (1992). An automated method for rotational correction and centering of three-dimensional functional brain images. *Journal of Nuclear Medicine*, 33, 1579-1585.
- Minoshima, S., Koeppe, R. A., Mintun, M. A., Berger, K. L., Taylor, S. F., Frey, K. A., & Kuhl, D. E. (1993). Automated detection of the intercommisural line for stereotactic localization of functional brain images. *The Journal of Nuclear Medicine*, 34, 322-329.
- Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.
- Pardo, J. V., Pardo, P. J., Janer, K. W., & Raichle, M. E. (1990). The anterior cingulate cortex mediates processing selection in the Stroop attentional conflict paradigm. *Proceedings of the National Academy of Sciences, USA, Neurobiology*, 87, 256-259.
- Petrides, M., Alivisatos, B., Evans, A. C., & Meyer, E. (1993). Dissociation of human mid-dorsolateral from posterior dorsolateral frontal cortex in memory processing. *Proceedings of the National Academy of Sciences, USA, Neurobiology*, 90, 873-877.
- Posner, M. I., Peterson, S. E., Fox, P. T., & Raichle, M. E. (1988). Localization of cognitive functions in the human brain. *Science*, 240, 1627-1631.
- Pylyshyn, Z. W. (1981). The imagery debate: Analogue media versus tacit knowledge. *Psychological Review*, 88, 16-45.
- Shallice, T. (1988). *From neuropsychology to mental structure*. Cambridge: Cambridge University Press.
- Smith, E. E., & Jonides, J. (1994). Working memory in humans: Neuropsychological evidence. In M. Gazzaniga (Ed.), *The cognitive sciences* (pp. 1009-1020). Cambridge, MA: MIT Press.
- Smith, E. E., Jonides, J., & Koeppe, R. A. (1996). Dissociating verbal and spatial working memory using PET. *Cerebral Cortex*. In press.
- Sternberg, S. (1969). The discovery of processing stages: Extensions of Donders' method. *Attention and Performance II. Acta Psychologica*, 30, 276-315.
- Talairach, J., & Tournoux, P. (1988). *A co-planar stereotaxic atlas of a human brain*. Stuttgart, Germany: Thieme.
- Tresch, M. C., Sinnamon, H. M., & Seamon, J. G. (1993). Double dissociation of spatial and object visual memory: Evidence from selective interference in intact human subjects. *Neuropsychologia*, 31, 211-219.
- Vanderplas, J. M., & Garvin, E. A. (1959). The association value of random shapes. *Journal of Experimental Psychology*, 57, 147-154.
- Warrington, E. K., & Rabin, P. (1970). Perceptual matching in patients with cerebral lesions. *Neuropsychologia*, 8, 475-487.
- Wilson, F. A. W., O'Scalaidhe, S. P., & Goldman-Rakic, P. S. (1993). Dissociation of object and spatial processing domains in primate prefrontal cortex. *Science*, 260:1955-1958.
- Worsley, K. J., Evans, A. C., Marrett, S., & Neelin, P. (1992). A three-dimensional statistical analysis for CBF activation studies in human brain. *Journal of Cerebral Blood Flow and Metabolism*, 12, 900-918.